

CLAIMS

What is claimed is:

1. A method for determining the memory state of a streaming media server, the method comprising:

identifying a size of a streaming media server's memory; and

determining a time interval that includes at least one access by at least one client to an amount of unique bytes of at least one streaming media file that are equal to the size of the streaming media server's memory.

2. The method of claim 1 further comprising:

for each of said at least one streaming media file accessed by said file accesses, determining at least one unique segment thereof included in said unique bytes.

3. The method of claim 2 further comprising:

associating with each of said at least one unique segment a corresponding timestamp of a most recent access thereof.

4. The method of claim 2 further comprising:

from the determined at least one unique segment, determining content of said streaming media server's memory.

5. The method of claim 2 further comprising:

receiving a request for a streaming media file to be served from the streaming media server; and

from the determined at least one unique segment, determining whether at least a portion of the requested streaming media file is in the streaming media server's memory.

6. The method of claim 5 wherein said determining whether at least a portion of the requested streaming media file is in the streaming media server's memory comprises:

determining, from the determined at least one unique segment, whether at least a first portion of the requested streaming media file is available in the streaming media server's memory such that the streaming media server is capable of immediately beginning to serve said at least a first portion from its memory in servicing the received request.

7. A system comprising:

server comprising a memory having a capacity, said server operable to serve at least one streaming file to clients communicatively coupled thereto; and

controller operable to determine an amount of unique bytes of said at least one streaming file served by the server during an elapsed time interval, wherein the amount of unique bytes are equal to the capacity of the server's memory.

8. The system of claim 7 wherein said controller is operable to determine said elapsed time interval.

9. The system of claim 8 wherein said controller is operable to determine said elapsed time interval responsive to receiving a new request from a client for access to a streaming file available from the server.

10. The system of claim 8 wherein said server serves said amount of unique bytes responsive to requests received from said clients during said elapsed time interval.

11. The system of claim 7 wherein said controller is operable to determine said amount of unique bytes responsive to receiving a new request for access to a streaming file available from the server.

12. The system of claim 11 wherein said unique bytes comprise most recently accessed unique bytes when said new request is received.

13. The system of claim 7 wherein said controller is operable to create a data structure that associates said unique bytes with a corresponding timestamp of a most recent access thereof.

14. The system of claim 7 wherein said controller is operable to determine unique segments of said at least one streaming file served during said elapsed time interval, wherein said unique segments include said unique bytes.

15. The system of claim 14 wherein said controller is further operable to determine a corresponding timestamp of a most recent access of each of said unique segments.

16. The system of claim 7 further comprising:
said controller is operable to determine, from the determined unique bytes of said at least one streaming file, content of said server's memory.

17. The system of claim 7 further comprising:
said controller is operable to receive a request for a streaming file to be served from the server and determine, from said unique bytes, whether at least a portion of the requested streaming file is in the server's memory.

18. A method for modeling the memory of a media server, the method comprising:
for each streaming file to which a media server provides access for its clients, determining at least one unique file segment, wherein each of said at least one unique file segment corresponds to a most recent access of a segment of the streaming file by a client of the media server;
identifying a memory management scheme implemented for the media server's memory;
identifying a maximum amount of bytes capable of being stored by the media server's memory; and
based at least in part on said memory management scheme and said maximum amount of bytes, determining from the at least one unique file segment content of the media server's memory.

19. The method of claim 18 wherein said memory management scheme is a Least Recently Used (LRU) management scheme.

20. The method of claim 18 wherein each of said at least one unique file segment comprises a sequential portion of a corresponding streaming file.

21. The method of claim 18 further comprising:
associating with each of said at least one unique segment a corresponding timestamp of a most recent access thereof.

22. The method of claim 18 wherein said determining content of the media server's memory further comprises:

determining the most recently accessed unique segments having a total number of bytes equaling the identified maximum amount of bytes capable of being stored by the media server's memory.

23. The method of claim 18 further comprising:

receiving a request for a streaming file to be served from the media server; and

from the determined content of the media server's memory, determining whether at least a portion of the requested streaming file is in the media server's memory.

24. Computer-executable software code stored to a computer-readable medium, the computer-executable software code comprising:

code for creating a segment-based data structure modeling the memory of a media server, wherein the data structure comprises identification of unique segments of streaming files previously accessed by clients of the media server with corresponding timestamps of most recent accesses of each unique segment.

25. The computer-executable software code of claim 24 wherein said unique segments of streaming files identify unique bytes of said streaming files accessed by clients of the media server that total an amount of bytes capable of being stored to the memory of said media server.

26. The computer-executable software code of claim 25 further comprising:

code for receiving a request for a streaming file to be served from the media server; and

code for determining, from the segment-based data structure, whether at least a portion of the requested streaming file is in the media server's memory.

27. The computer-executable software code of claim 26 further comprising:

code for determining whether a received request is capable of being serviced by the media server without resulting in an overload of the media server.

28. A method for modeling the memory of a streaming media server, the method comprising:

creating a segment-based access model for at least one streaming media file available on a streaming media server; and

using said segment-based access model to construct a segment-based model of the streaming media server's memory.

29. The method of claim 28 wherein said creating a segment-based access model comprises creating said segment-based access model for a time interval of interest.

30. The method of claim 29 wherein said creating a segment-based access model comprises:

determining, for said time interval of interest, at least one unique segment of each of said at least one streaming media files accessed during said time interval of interest by at least one client of said streaming media server.

31. The method of claim 30 wherein said determining said at least one unique segment comprises:

determining, for each of said at least one streaming media files, if multiple accesses were made to the streaming media file during said time interval of interest; and

if multiple accesses were made to the streaming media file during said time interval of interest, then determining from the multiple accesses said at least one unique segment of the streaming media file.

32. The method of claim 29 wherein said segment-based model of the media server's memory represents said media server's memory as of an end of the time interval of interest.

33. The method of claim 32 wherein said end of the time interval of interest comprises a time at which a new request for access to said at least one streaming media file is received.

34. The method of claim 28 wherein said segment-based access model comprises an identification of at least one unique segment of a streaming media file with a corresponding identification of a time of a most recent access of said at least one unique segment.

35. The method of claim 28 wherein said segment-based model of the streaming media server's memory comprises (a) identification of unique segments of streaming media files previously accessed by clients of the streaming media server and (b) identification of corresponding timestamps of most recent accesses of each unique segment.

36. The method of claim 35 wherein said unique segments of streaming media files identify unique bytes of said streaming media files accessed by clients of the streaming media server that total an amount of bytes capable of being stored to the streaming media server's memory.

37. The method of claim 28 further comprising:
receiving a request for a streaming media file to be served from the streaming media server; and

from the segment-based model of the streaming media server's memory, determining whether at least a portion of the requested streaming media file is in the streaming media server's memory.